IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

| | | |
|---|---|---|
| EMC CORPORATION AND EMC ISRAEL DEVELOPMENT CENTER, LTD., | ) | |
| Plaintiffs, | ) | |
| v. | ) | C.A. No. 12-956 (GMS) |
| ZERTO, INC., | ) | |
| Defendant. | ) | |

**PLAINTIFFS EMC CORPORATION AND EMC ISRAEL DEVELOPMENT CENTER, LTD.'S OPENING CLAIM CONSTRUCTION BRIEF**

Plaintiffs EMC Corp. and EMC Israel Development Center, Ltd. (collectively "EMC"), by and through their counsel, and pursuant to this Court's Revised Scheduling Order, provide their Opening Claim Construction Brief regarding the construction of disputed claim terms in U.S. Patent Nos. 6,073,222; 7,577,867; 7,603,395; 7,647,460; and 7,971,091.

MORRIS, NICHOLS, ARSHT & TUNNELL LLP
Jack B. Blumenfeld (#1014)
Michael J. Flynn (#5333)
1201 North Market Street
P.O. Box 1347
Wilmington, DE 19899
(302) 658-9200
jblumenfeld@mnat.com
mflynn@mnat.com

*Attorneys for Plaintiffs*

OF COUNSEL:

Michael N. Rader
Gregory F. Corbett
Allen S. Rugg
Nathan R. Speed
Elizabeth A. DiMarco
WOLF, GREENFIELD & SACKS, P.C.
600 Atlantic Avenue
Boston, MA 02210-2206
(617) 646-8000

Paul T. Dacier
Krishnendu Gupta
William R. Clark
EMC CORPORATION
176 South Street
Hopkinton, MA 01748

April 11, 2014

# TABLE OF CONTENTS

**I. U.S. Patent No. 7,647,460 ..... 1**
A. "storage system" ..... 2
B. The "mirror" Terms ..... 3
C. "means . . . for mirroring" ..... 5
D. "controller" ..... 6

**II. U.S. Patent No. 6,073,222 ..... 8**
A. "preservation memory" ..... 8
B. "when a new block written to the mass storage system is to overwrite the original data block" ..... 9
C. "original data block" ..... 10

**III. U.S. Patent Nos. 7,603,395 and 7,971,091 ("the Bingham Patents") ..... 10**
A. "pseudosnapshot" ..... 12
B. The "surfacing" Terms ..... 14
C. "event marker" ..... 15
D. The "quiesce" Terms ..... 16
E. "write journal" ..... 18
F. "client file index" ..... 19

**IV. U.S. Patent No. 7,577,867 ..... 19**
A. "data protection appliance (DPA) manager" ..... 20

**V. CONCLUSION ..... 20**

# TABLE OF AUTHORITIES

**Cases**

*Aventis Pharma S.A. v. Hospira, Inc.*,
675 F.3d 1324 (Fed. Cir. 2012).....2

*Home Diagnostics, Inc. v. LifeScan, Inc.*,
381 F.3d 1352 (Fed. Cir. 2004).....2

*Phillips v. AWH Corp.*,
415 F.3d 1303, 1314 (Fed. Cir. 2005).....4

*Symantec Corp. v. Computer Associates International, Inc.*,
522 F.3d 1279 (Fed. Cir. 2008).....3

*Thorner v. Sony Computer Ent. Am.*,
669 F.3d 1362 (Fed. Cir. 2012).....2

*XpertUniverse, Inc. v. Cisco Sys., Inc.*,
09-cv-157-RGA, 2012 WL 1416776 (D. Del. Apr. 20, 2012).....7, 12

**Statutes**

35 U.S.C. § 112.....6

## TABLE OF EXHIBITS

The following exhibit cited herein refers to the exhibits to the Declaration of Nathan R. Speed In Support of Plaintiffs' Proposed Claim Constructions, filed concurrently herewith. All other citations refer to the exhibits to the parties' Joint Appendix of Intrinsic and Dictionary Evidence (D.I. 94).

| EXHIBIT | DESCRIPTION |
|---|---|
| 3 | U.S. Patent No. 5,935,260 |

This litigation and the five asserted patents concern data protection. Companies today generate a staggering amount of data. This data is the lifeblood of the global economy and its protection is not simply important, it is essential. The five asserted patents disclose innovative techniques for both protecting data and recovering the protected data if needed. These patents and the products they protect have helped make EMC a global leader in data storage and management.

**I. U.S. Patent No. 7,647,460**

The '460 patent discloses a data protection system in which two storage systems are connected via at least one communication link. (JA84, Abstract; JA86, Fig. 2.) As data is written to the source storage system, it is transferred over the link to a target storage system where it is also written. (JA90, 1:40-42.) This process results in the target storage system maintaining a continually updated mirror image of the source storage system, and provides an alternate source of data in the event that the source storage system fails or is destroyed, e.g., by a fire, natural disaster, etc. The prosecution history reveals that the '460 patent improved on conventional "remote mirroring" systems. In a first conventional system, the host device that issued a write (i.e., a request to store data) to the source storage system also issued a second write over a communication link to the target storage system. In a second conventional system, the host was unburdened from having to issue two writes, as mirroring controllers were added to the storage systems and those controllers communicated directly to mirror writes made to the source storage system to the target. In this system in which mirroring was handled directly by the storage systems, a dedicated communication line was used to connect the two storage systems. (*Id.*, 1:62-2:21.) The '460 patent recognized that dedicated lines were expensive and could only be implemented over relatively short distances. (*Id.*, 2:14-30.) The patent describes using a non-dedicated network cloud to connect the storage systems. This addressed two limitations of the prior art because the cloud could be implemented on a low-cost public network (e.g., the Internet), and data could be transmitted over that network to any remote location without geographic constraint. (JA90-91, 2:62-3:21.)

**A. "storage system"**

| | |
|---|---|
| **EMC**: "the set of components that stores and controls the storage of information written from the CPU, including one or more storage devices and one or more controllers" | **Zerto:** "the combination of a storage device and its storage controller" |

The parties agree that a "storage system" includes at least one controller and at least one storage device. The dispute is whether a storage system may include more than those two components, as EMC proposes, or whether it is limited to ***only*** one storage device and one controller dedicated exclusively to controlling that storage device, as Zerto proposes.

Absent an express definition or claim-scope disavowal, "storage system" should be construed according to its customary meaning. *Aventis Pharma S.A. v. Hospira, Inc.*, 675 F.3d 1324, 1331 (Fed. Cir. 2012) ("[W]e will only interpret a claim term more narrowly than its ordinary meaning under two circumstances: 1) when a patentee sets out a definition and acts as [its] own lexicographer, or 2) when the patentee disavows the full scope of a claim term either in the specification or during prosecution." (internal quotation omitted).) The IEEE Technical Committee on Mass Storage Systems and Technology defines "storage system" as: "a set of devices, associated software, operators, physical processes, etc., that provides services for the storage and retrieval of data." (JA209.) This definition is consistent with the Dictionary of Computer Terms which defines "system" as: "an organized collection of components that have been optimized to work together in a functional whole." (JA241.) Thus, the customary meaning of "storage system," in the context of claim 1, is, as EMC proposes, the set of components that work together to store information written from the CPU.

The two limited exceptions that permit a court to interpret a term more narrowly than its customary meaning do not apply. First, "storage system" is not expressly defined in the specification. Patentees must "clearly express an intent" to define a term and act as a lexicographer. *Thorner v. Sony Computer Ent. Am.*, 669 F.3d 1362, 1365 (Fed. Cir. 2012). The inventors of the '460 patent expressed no such intent. Second, neither the specification nor the prosecution history demonstrates a "clear disavowal" of the full scope of "storage system." *Home Diagnostics, Inc. v. LifeScan, Inc.*, 381 F.3d 1352, 1358 (Fed. Cir. 2004) ("Absent a clear

disavowal … in the specification or the prosecution history, the patentee is entitled to the full scope of its claim language.").

Contrary to its customary meaning, Zerto's construction limits "storage system" to a particular configuration: a storage device and "its" controller. This construction is inconsistent not only with the term's customary meaning, but with several embodiments described in the '460 patent. For example, Figs. 2 and 4 disclose source storage systems that each includes ***multiple*** controllers 5 and 19, and target storage systems that each includes ***multiple*** controllers 7 and 19. Additionally, the patent identifies EMC's SYMMETRIX line of disc arrays as one example of a storage device (3 and 9 in Fig. 2), and the SYMMETRIX product is well known to include ***multiple*** different controllers. (JA90, 1:23-25; *see also* Ex. 3 at 1:24-46 (describing the SYMMETRIX product as including "several types of controllers for controlling the various aspects of the data transfers associated with the storage system").) Zerto's proposal is incompatible with the patent's description of storage systems that include multiple controllers.

Zerto is not the first accused infringer to contend that "system" should be narrowly construed. In *Symantec Corp. v. Computer Associates International, Inc.*, the defendant argued that "computer system" should be limited to a single computer because the specification only disclosed embodiments in which a single computer was used. 522 F.3d 1279, 1290 (Fed. Cir. 2008). The district court adopted the defendant's construction, but the Federal Circuit reversed, holding that absent a special definition in the patent, the customary meaning of "system," i.e., one or more related components, controlled. *Id.*

Pursuant to both its customary meaning and the specification, "storage system" may include, along with at least one storage device and at least one controller, additional components that serve a common storage purpose. EMC's proposed construction should therefore be adopted.

**B. The "mirror" Terms**[1]

[1] EMC disagrees with Zerto's attempt to construe "mirror" in isolation and provides a construction for the phrases in which it appears. *Phillips v. AWH Corp.*, 415 F.3d 1303, 1314 (Fed. Cir. 2005) ("[T]he context in which a term is used in the asserted claim can be highly instructive.").

| | | |
|---|---|---|
| 1. "mirror" | **EMC**: N/A | **Zerto**: "create an identical copy of data that exists on one system on a second system" |
| 2. "to mirror at least some of the information written from the CPU to the first storage system in the second storage system" | **EMC**: "maintain a continually updated copy of data that exists on a portion of the first storage system on the second storage system by copying each write to the portion of the first storage system to the second storage system" | **Zerto**: N/A |
| 3. "to mirror at least some of the information written from the CPU to the first storage system in both the second and third storage systems" | **EMC**: "maintain a continually updated copy of data that exists on a portion of the first storage system on the second and third storage systems by copying each write to the portion of the first storage system to the second and third storage system" | **Zerto**: N/A |

Zerto's construction is inconsistent with the ordinary meaning of "mirror" and the intrinsic record. "Mirroring" is a technical term with a particular meaning. This meaning, which is also the only meaning consistent with the intrinsic record, is "to maintain a continually updated copy of data by copying each write written to a first storage system to a second storage system." Thus, like a household mirror, the data on the second storage system continually reflects or mirrors the data on the first storage system. Zerto's proposal encompasses a wide-range of copying techniques that are not mirroring, including backup techniques that the applicants distinguished from "mirroring" during prosecution.

The specification describes mirroring as follows: "As data is written to the source storage device 3, it can also be written and mirrored to the target storage device 9." (JA90, 1:39-41.) The specification describes two embodiments. In one, a host CPU cannot issue a new write until it receives confirmation that its prior write was stored not only on the local storage system, but also "on the target storage system." (JA95, 11:44-66.) In the other, a "store and forward" technique is used wherein the host CPU is released when data is stored not only on the source storage system, but also in a temporary storage area. The data is "subsequently transferred from the temporary storage location to the target storage system." (*Id.*, 11:47-12:3.) These descriptions make clear that "mirroring" requires storing to a second storage system every write that a CPU issues to a mirrored portion of the first storage system.

The claim language reinforces this understanding of the term. The claims require a "mirroring controller" (or "controller") to mirror information "in response to" or "responsive to" "the information being written from the CPU to the first storage system." The "responsive" language was added during prosecution to clarify that "mirroring" occurs as data is written to a first storage system, and not later in response to a backup utility being run. (JA166 ("Ohran is not directed to a system that performs a mirroring of data. Rather, ... Ohran is directed to a backup system that transfers data that has previously been stored in a mass storage device to a backup system…The backup process…is therefore not responsive to data being written to the primary mass storage device.").) The remainder of the prosecution history also supports EMC's proposal. For example, the applicants argued that a prior art reference disclosing a "backup system" did not teach mirroring because "mirroring updates are made in essentially real time," while a "backup operation is typically performed off-line." (JA144.) The applicants were clear: mirroring is a specific technique for copying data that involves continually storing at a second storage system writes that a CPU issues to a first storage system so that the second storage system "mirrors" the first.

Zerto's proposal does not explain to the jury how an identical copy of data is created on a second storage system, in an apparent attempt to include techniques that are not mirroring. For example, a backup system that copies data existing on a first storage system to a second storage system once a day would satisfy Zerto's proposed construction, yet, as noted above, the applicants went to considerable lengths to distinguish such backup techniques from "mirroring" during prosecution. Zerto's construction is inconsistent with both the term's ordinary meaning and the intrinsic evidence, and should be rejected.

**C. "means . . . for mirroring"**

| | |
|---|---|
| **EMC**: Function: mirroring at least some of the information written from the CPU to the first storage system in the second storage system<br>Structure: source and target storage controllers, network RDF controller, and network interface unit | **Zerto:** Function: mirroring at least some of the information written from the CPU to the first storage system in the second storage system<br>Structure: source and target storage controllers, network RDF controller, network interface unit, and a network cloud |

Despite repeated requests for clarification, Zerto waited until the day before the opening claim construction briefs were due to clarify that it agreed with EMC that the "means … for mirroring" limitation is governed by 35 U.S.C. § 112, ¶ 6, and waited until the day the briefs were due to revise its proposed construction to agree largely with EMC's proposal.

The sole dispute remaining between the parties is whether a "network cloud" is part of the "means … for mirroring" limitation. The claim language dictates that it is not. The "means … for mirroring" limitation is found in dependent claims 11 and 13 (which depend from independent claim 1) and dependent claim 46 (which depends from independent claim 44). Each of the dependent claims requires that the "mirroring controller" or "controller" of its independent claim include "means … for mirroring." Independent claims 1 and 44 each recites a system with separately recited elements of a "mirroring controller" and at least one "communication link" that includes a "network cloud" (or one of an intranet or the Internet). The "means … for mirroring" is specifically recited as further limiting the "mirroring controller." Requiring a "network cloud" to be part of the "mirroring controller" is incompatible with the claims' treatment of the terms as separate elements. Furthermore, the specification nowhere links the function of the "mirroring controller" to the Internet or any other network cloud; it defies logic to require that the controller somehow include a "network cloud" such as the Internet. Zerto's proposal should be rejected.

**D. "controller"**

| | |
|---|---|
| **EMC**: "a component of the storage system that controls at least one function of the storage system" | **Zerto:** "a device that manages the operation of and connection to another device" |

Zerto asserts that a controller is a "device" that manages the "operation of and connection to another device." This construction is unclear. It does not explain to the jury what a "device" is, does not identify the "another device" whose operation is purportedly "managed," and does not identify the third component whose "connection to" the "another device" the controller purportedly manages. These ambiguities alone are reason to reject Zerto's proposal. *XpertUniverse, Inc. v. Cisco Sys., Inc.*, 09-cv-157-RGA, 2012 WL 1416776, at *3 (D. Del. Apr.

20, 2012) ("There is no need to inject further language and potential confusion into the claims when, as here, the claims and context make [a term's meaning] clear.").

To the extent Zerto's proposal can be understood, the parties' constructions appear to differ both on what a controller is and what a controller does.

EMC proposes that the controller is a "component of the storage system," while Zerto proposes that it is a "device." Zerto's overly narrow proposal finds no support in the patent, which never describes a controller as a "device." The only "devices" identified in the patent are storage devices 3 and 5, which are not controllers. Zerto's proposal appears intended to limit a controller to a stand-alone device. That is inconsistent with the specification, which explains, for example, that both storage controllers and network RDF controllers can be implemented in numerous ways, including as software executed by any processor. (JA93, 7:14-21; 7:33-35.)

With respect to the function performed by a controller, each of the claims in which the term "controller" appears already explicitly recites that controller's function. Given that Zerto sought to construe "controller" out of context, it is inappropriate to impose any particular function onto the generic term "controller" because as the intrinsic evidence make clear, different types of controllers do different things. This is why EMC proposes simply explaining to the jury that a controller controls at least one function, but letting the function explicitly recited in each claim speak for itself. Zerto's construction should be rejected because the jury will likely be confused about how Zerto's "manages the operation of and connection to another device" functional language relates to the functions already explicitly recited in the claims.

In addition, Zerto's construction narrows the claims inappropriately. None of the claimed functions requires a controller to "manage[] the operation of and connection to another device," and the parties have already agreed that a "mirroring controller" is "one or more controllers configured to mirror." The patent nowhere describes a "controller" as performing the "managing and connection to" functions that Zerto proposes. For example, the patent explains that the storage controllers "perform the function of implementing the protocol for transferring information between them" (JA93, 7:5-8), and the network RDF controllers format data for

transmission over a network cloud (*id.*, 7:22-39). Zerto's proposal requires a controller to perform functions that are not described in the patent, are not the functions the claims explicitly recite, and are not the functions that the parties agreed upon for a "mirroring controller." Zerto's proposal should be rejected.

## II. U.S. Patent No. 6,073,222

The '222 patent discloses a system and a method for (a) preserving data as it existed on a mass storage system at a particular moment in time and (b) accessing the preserved data as it existed at that previous point in time, even if the data was subsequently overwritten. (JA1, Abstract; JA245, Fig. 2[2].) To do this, when a data block on a mass storage system is overwritten, the block is preserved in a "preservation memory." (JA29, 31:10-15; JA245 (Fig. 2); JA251, 5:48-60.) A "virtual device" is created that can access data in the preservation memory and the mass storage system. This provides a user with access to data as it existed at a particular moment in time, even if the accessed data includes some data blocks located on the mass storage system (because they have not been overwritten) and other data blocks preserved in the preservation memory (because they have been overwritten). (JA1, Abstract; JA251, 5:66-6:10.) By preserving each data block only when it is overwritten and establishing a virtual device that provides access to data blocks in both the preservation memory and the mass storage system, the invention consumes considerably less storage resources than prior art techniques for preserving and accessing data as it existed at a particular moment in time. (JA249-50, 2:65-3:5.)

### A. "preservation memory"

| | |
|---|---|
| **EMC**: "storage used for preserving data blocks" | **Zerto**: "memory dedicated to holding a snapshot of a portion of memory at the selected moment" |

EMC's proposed construction for "preservation memory" explains the meaning of this straightforward limitation, while Zerto's proposed construction complicates and improperly narrows the term. The specification repeatedly explains that the "preservation memory" is used to store, and thereby preserve, data blocks that are overwritten. (JA1, Abstract; JA18, 9:1-12; JA22,

[2] The '222 patent incorporates by reference related U.S. Patent 5,649,152. (JA14, 1:6-23.)

17:47-18:2.) This is consistent with the claims. For instance, claim 1 recites that an original data block is preserved "by placing the original data block in the preservation memory." (JA29, 31:10-14.) The specification clarifies that the preservation memory "may comprise any type of writable storage device … and may optionally be a portion of [the] mass storage device." (JA18, 9:5-11.)

Despite the patent's broad description of a "preservation memory," Zerto seeks to narrow the term in two ways, by asserting that it be "***dedicated*** to holding a ***snapshot*** of a portion of memory at the selected moment." The intrinsic evidence does not support Zerto's narrowing construction. The patent provides no reason for the "preservation memory" to be "dedicated" to any function. To the contrary, it specifically states that the preservation memory may be part of the mass storage device which serves other functions. (JA18, 9:5-11.)

Moreover, Zerto's proposal injects the concept of a "snapshot" into the asserted claims where it does not belong. The specification describes creation of a "snapshot" as but one embodiment (e.g., JA18, 19:7-9) and makes clear that the claims are not limited to the illustrative embodiments (JA28, 30:57-65). Moreover, the patent includes a number of claims that are unasserted and, unlike the asserted claims, recite the snapshot concept. (*Compare* asserted claim 1 ("… a method of accessing data as it existed…" *with* unasserted claim 11 ("… a method for accessing ***a snapshot*** of the mass storage system...").) As discussed in Section III. A. below, "snapshot" is a technical concept with a particular meaning. The asserted claims, which do not recite "snapshot," are not limited to snapshots, and Zerto's attempt to inject that concept into the asserted claims impermissibly limits their scope and should be rejected.

**B. "when a new block written to the mass storage system is to overwrite the original data block"**

| | |
|---|---|
| **EMC**: Plain and ordinary meaning. If the Court further construes this term, EMC proposes the following as its plain and ordinary meaning: "*in response to a new block written to the mass storage system that is to overwrite the original data block*" | **Zerto**: "before a new block written to the mass storage system is to overwrite the original data block" |

This phrase is used in claims 1 and 20. In the context of those claims, the phrase is readily understood and needs no construction. However, if the Court determines that a construction is

necessary, it should adopt the phrase's customary meaning: "***in response to*** a new block written to the mass storage system that is to overwrite the original data block."

Claims 1 and 20 require an "original data block" to be placed in the preservation memory "***when*** a new block written to the mass storage system is to overwrite the original data block." A jury will understand what the phrase means in the context of the claims. Yet, contrary to the clear language, Zerto seeks to narrow the claim by injecting a limitation that requires an ***order*** of operations by replacing "when" with "before," thereby requiring the "original data block" to be placed in the preservation memory ***before*** a new block is written to the mass storage to overwrite the original data block. This order of operations does not appear in the claims and is contrary to the customary meaning of the word "when," which encompasses "***after***" – the exact opposite order that Zerto's unduly narrowing construction imposes. (JA260 (defining "when" as "at a time at which; at the time at which; just as; ***after***: *I found it easily when I started to look seriously*").) Zerto's overly narrow construction should be rejected.

**C. "original data block"**

| | |
|---|---|
| **EMC**: Plain and ordinary meaning. | **Zerto**: "a block of data stored on the mass storage system at the time the snapshot was taken" |

In the context of claims 1 and 20, the term "original data block" is readily understood and needs no construction because the claim explicitly recites "an original data block as it existed in the mass storage system at the selected moment." *U.S. Surgical Corp. v. Ethicon, Inc.*, 103 F.3d 1554, 1568 (Fed. Cir. 2007) ("[Claim construction] is not an obligatory exercise in redundancy."). Moreover, Zerto's proposal should be rejected because it attempts to introduce the "snapshot" concept that is recited in other claims but not the asserted ones.

## III. U.S. Patent Nos. 7,603,395 and 7,971,091 ("the Bingham Patents")

The Bingham patents share a common specification.[3] The patents disclose novel systems and methods for performing replication operations on a continuous data protection ("CDP") system. (JA73, 1:20-23.)

[3] Unless otherwise noted, all citations are to the specification of the '395 patent.

Write transactions that an application issues to a primary storage system are also stored in a separate location, referred to in the Bingham patents as a "CDP engine." (JA73, 2:19-30; JA67, Fig. 1.) By separately storing a copy of each write transaction, "the CDP engine is able to retain information about every change to a selected application, directory, volume, or the like over a period of time" (JA73, 2:26-28), and the user can use the CDP engine to access stored data as it appeared at a given point in time in the past (*id.*, 2:29-30). Thus, CDP permits recovery of data at any prior point-in-time.

"Traditional snapshot" technology creates "a read-only point-in-time replica or representation" of data. (JA74, 3:46-48.) Using traditional snapshot techniques requires creating a distinct replica or representation of data every time a snapshot was taken, so that taking snapshots at regular time intervals would "consume [significant] storage space and resources." (*Id.*, 3:52-53.)

As the specification and prosecution history make clear, the Bingham patents do not use snapshots – they use "pseudosnapshots." Generating pseudosnapshots only requires inserting "event markers" into the CDP engine at desired points in time, thereby creating "earmarked datasets within the CDP engine that can later be reliably surfaced as a copy." (*Id.*, 3:30-35.) Any pseudosnapshot can later be recovered by reading from the CDP engine the "earmarked" data that was stored at the time associated with the event marker. Data is only "surfaced" for specific pseudosnapshots that are needed. (*Id.*, 3:35-40.) Thus, while every pseudosnapshot is available to the user as an option to recover, the system consumes far less storage than would be required to store actual snapshots at every event marker. (*Id.*, 3:40-45.)

In an illustrative embodiment, each time an application issues a write transaction, a copy of the transaction is stored in a write journal on the CDP system. (JA73, 2:19-30; claim 1, '395 patent.) A replication application creates "event markers" that identify points in time in the series of write transactions stored in the write journal. (JA79, 13:60-65; JA80, 15:28-67; JA71, Fig. 5.) The replication application may use an event marker to "surface" a copy of the data corresponding to the write transactions the event marker identifies. (JA80, 16:3-17; *id.*, 16:43-47;

JA71, Fig. 5.) An application can be quiesced (i.e., instructed to suspend writes) before the event marker is inserted, which ensures that the data that can be recovered for that event marker is in an "application consistent state." (JA79, 13:39-42; JA81, 18:27-42.) Once surfaced, the data corresponding to the pseudosnapshot can be replicated by being copied to a storage device or having another replication operation performed on it. (JA74, 4:6-8.)

**A. "pseudosnapshot"**

| | |
|---|---|
| **EMC**: "a dataset, identified by an event marker, within the continuous data protection system that can be surfaced" | **Zerto**: "an item, presented to a user as a traditional snapshot and associated with an event marker, that can be used to surface a copy of application data as it existed at the time of the event marker" |

The term "pseudosnapshot" has no customary meaning – it was coined by the inventors. Therefore, the patents' description of the term is controlling. The patents describe "pseudosnapshots" as "earmarked datasets within the CDP engine that can later be reliably surfaced as a copy." (JA74, 3:33-35.) The patents further explain that the datasets are earmarked using event markers. (JA79, 14:7-12 ("[W]hen a pseudosnapshot is created, no snapshot is actually created on any storage disk. Instead, an event marker is created at the time the pseudosnapshot is created to allow a copy of the data to be surfaced from the CDP engine 110 up until the point in time at which the event marker was created."); *Id.*, 14:46-47 ("Graphical user interface 122 allows the user to manage all the event markers.").). EMC's proposed construction aligns directly with the patents' description of a pseudosnapshot.

Zerto's proposed language is inconsistent with the intrinsic evidence and injects ambiguity into the claims. It is not at all clear what it means to "present" an "item" "as a traditional snapshot." To the extent that Zerto's proposed construction implies that a "pseudosnapshot" must have some ***unspecified*** properties of a traditional snapshot, the language is impermissibly ambiguous and should be rejected. *XpertUniverse, Inc.*, 2012 WL 1416776 at *3.

Moreover, the language "presented to a user as a traditional snapshot" is inconsistent with the intrinsic evidence and appears to be another attempt by Zerto to inject the concept of a "snapshot" where it does not belong. Like "mirror" discussed above, "snapshot" is a technical

term that carries special meaning. The patents explain that a "snapshot" is a "read-only point in time ***replica or representation*** of the original data or contents of one or more volumes" that "consume[s] storage space and resources." (JA74, 3:46-53.) This is consistent with its customary meaning. A snapshot is a particular type of data structure. As the patent indicates, there are two main types of "snapshots," i.e., a "replica" and a "representation." (*Id.*, 3:47.) A "split-mirror" is a replica type and creates a copy "of the entire volume." (JA274.) A "copy-on-write" is a representation type and creates a copy "of changes to stored data every time new data is entered or existing data is updated." (*Id.*) The term "snapshot" connotes a particular set of technologies not used in the Bingham patents, and Zerto's attempt to inject those technologies into "pseudosnapshots" is untenable in light of the intrinsic evidence.

The patents are clear that, as its very name indicates, a "pseudosnapshot" is ***not*** a snapshot. Unlike a snapshot, a "pseudosnapshot" is not a "replica or representation" of data when it is created. (JA74, 3:40-64.) This distinction allows a user to create millions of pseudosnapshots without consuming the resources that traditional snapshots consume. (*Id.*; JA80, 15:12-16.) This distinction was emphasized throughout prosecution of the patents. (JA184 ("Advantageously,… pseudosnapshots do not require near the resources of a traditional snapshot. As a result, numerous pseudosnapshots can be saved using event markers – thus reducing the amount of resources and storage that would be required for traditional snapshots."); JA196 ("[A] pseudosnapshot is distinct from a conventional snapshot at least because the dataset being surfaced is not stored in the pseudosnapshot itself."); JA197 ("[A] pseudosnapshot identifies data in the continuous data protection system – the pseudosnapshot does not store the actual data itself.").)

Furthermore, to the extent Zerto's proposal implies that "an item" is not a "pseudosnapshot" until it is "presented to a user," Zerto's proposal excludes the vast majority of "pseudosnapshots" described in the specification. The patents are clear that "pseudosnapshots" are "generated in the CDP engine by creating event markers." (JA74, 3:30-35.) Once an event marker is created, the "pseudosnapshot" has been "generated" and is an earmarked dataset "that

can ***later*** be reliably surfaced as a copy." (*Id.*, 3:33-35.) Zerto's construction, however, appears to limit "pseudosnapshots" to "items" that are actually "presented to a user." That construction is inconsistent with the patents, which never link the presentation of a "pseudosnapshot" with its generation; indeed, the patents explicitly teach that the two concepts are distinct. For example, each independent claim of the '395 patent explicitly de-couples the generation of a "pseudosnapshot" from its presentation to a user. (*See, e.g.*, claim 9 ("generating a plurality of pseudosnapshots … presenting the plurality of pseudosnapshots to a user…").) Moreover, the patents explain that "millions" of "pseudosnapshots" can be generated (JA80, 15:8-16), yet never describe an embodiment in which "millions" of "pseudosnapshots" are "presented to a user."

Because Zerto's proposal is inconsistent with the intrinsic evidence, it should be rejected. EMC's proposal, which is fully consistent with the intrinsic record, should be adopted.

**B. The "surfacing" Terms**[4]

| | | |
|---|---|---|
| 1. "surfacing" | **EMC**: N/A | **Zerto**: "generating a snapshot of data as it existed at a given point in time from a series of transactions stored in a write journal" |

| | | |
|---|---|---|
| 2. "surfacing [a/the] copy of [a/the] at least a portion of the series of transactions" | **EMC**: "generating [a/the] data structure from the dataset of [a/the] at least a portion of the series of transactions" | **Zerto**: N/A |
| 3. "surfacing of a copy of transactions" | **EMC**: "generating a data structure from the dataset of a series of transactions" | **Zerto**: N/A |

In an illustrative embodiment of the Bingham patents, a copy of all write transactions from a production application are stored in a write journal on a CDP engine. (JA73, 2:19-30; JA75, 5:34-40.) The patents explain that "surfacing" is the act of generating a "data structure" from the data within the write journal upon which a user can perform various operations, referred to in the patents as "replication operations." (JA73, 2:11-15; JA75, 5:34-40.) Like components of an automobile on an assembly line that have not yet been assembled to form an automobile, the data in the write journal that is earmarked by an event marker is not accessible as a dataset unless and until a data structure is generated that "assembles" the data into an accessible dataset. The act of

[4] Zerto seeks to construe "surfacing" in isolation. The term should not be construed in isolation and EMC therefore provides a construction for the phrases in which the term appears.

generating this data structure that "assembles" the data into an accessible data structure is "surfacing."

The parties agree that "surfacing" involves "generating" something from the data of a series of transactions. The dispute is what the "something" is that is generated. EMC proposes, consistent with the intrinsic evidence, that the "something" is a "data structure." Zerto, on the other hand, proposes that surfacing generates a "***snapshot*** of data." The intrinsic evidence is clear that it does not.

As discussed above, "snapshot" is a technical term that carries special meaning and does not cover ***any*** point-in-time copy of a dataset. The patents explain, contrary to Zerto's proposal, that the "something" generated from "surfacing" is a "data structure," that is not a "snapshot." For example, when describing "exemplary methods for surfacing a copy using a pseudosnapshot" (JA80, 16:1-2), the patents state that surfacing involves formatting data within a CDP engine "into a ***data structure*** that has the ***appearance of a traditional snapshot***" (*id.*, 16:14-16). In other instances, the patents explain that a "surfaced copy ***can be*** formatted to have the ***properties of*** a conventional read-only snapshot" (JA79, 14:12-14) or "can be utilized ***similar*** to a traditional snapshot" (JA76,7:28). The patents never state that the generated data structure ***is*** a snapshot, and by contrasting the data structure with a snapshot, the patents make clear that it is ***not*** a snapshot.

The intrinsic evidence is clear -- surfacing generates a "data structure" that has some similarities to, ***but is not***, a snapshot. Because EMC's proposal is consistent with the intrinsic evidence and Zerto's is not, EMC's proposal should be adopted.

**C. "event marker"**

| | |
|---|---|
| **EMC**: "a data structure that includes information that can be used to identify a pseudosnapshot" | **Zerto**: "a data structure that includes the date and time at which the pseudosnapshot was created" |

The primary dispute over "event marker" is whether, as Zerto urges, it must include "the date and time at which the pseudosnapshot was created." The intrinsic record makes clear that it need not. The patents state that the "event marker is a data structure that ***may exemplarily*** contain

information regarding the time at which the pseudosnapshot was created" and that the claims are not limited to this (or any other) illustrative embodiment. (JA79, 13:56-60; JA82, 19:23-30.)

The patents also describe an embodiment wherein the event marker does ***not*** contain time and date information. (JA80, 15:60-65.) After explaining that the event markers need not be stored with the data in the write journal because they "***may merely*** include the time and date at which the pseudosnapshot is created," the patents describe an alternative embodiment -- "[a]lternatively, the first entry [in the write journal] immediately after an event marker is created could also be used to identify the existence of a pseudosnapshot." (*Id.*, 15:62-65.) The next entry after the event marker will have the time and date at which the next change to the data occurred after the event marker is established. Thus, the event marker's position in the journal indirectly provides the date and time of the pseudosnapshot, so this information need not be included in the event maker itself. Accordingly, the patent makes clear that although an event marker ***may*** include date and time information in some embodiments, it is not required to include that information and, in fact, in some embodiments it does not. Zerto's proposal is inconsistent with the intrinsic evidence and should be rejected.

EMC's reading of "event marker" is consistent with the specification and the claims, some of which (claim 1, '091 patent) expressly require that the event marker identify "a time and date" at which a pseudosnapshot was created, others of which do not (claim 1, '395 patent).

**D. The "quiesce" Terms**

| | | |
|---|---|---|
| 1. "quiesce" | **EMC**: "in response to an instruction to do so, enter an inactive state wherein the production application does not generate write operations" | **Zerto**: "pause write activity" |
| 2. "quiescing"[5] | **EMC**: "instructing the production application to enter an inactive state wherein the production application does not generate write operations" | **Zerto**: "pausing write activity" |
| 3. "quiescent state" | **EMC**: "a state that results from the production application executing an instruction to quiesce" | **Zerto**: "a state in which write activity is paused" |

[5] Zerto seeks to construe "quiescing" in isolation. The term should not be construed in isolation and EMC therefore provides a construction for the phrase "quiescing the production application."

These are terms of art that Zerto again seeks to construe inconsistent with their customary meaning and the intrinsic evidence. While the parties agree that a "quiesced" application is in an "inactive state" (JA79, 13:40-42), they dispute whether an application must be instructed to quiesce, as EMC proposes, or whether an application is "quiesced" any time there is a pause in its write activity, as Zerto proposes. As shown below, the intrinsic evidence and the term's ordinary meaning both support EMC's proposal and contradict Zerto's.

The Bingham patents describe the ability to recover data in an "application consistent state." (JA81, 18:27-42.) The parties agree that an "application consistent state" is "a point in time when all pending write operations for an application have been committed to disk." (D.I. 86 at 3.) To achieve this state, the Bingham patents teach taking the affirmative step of "quiescing" a production application prior to inserting an event marker in the write journal. (JA81, 18:32-34.)

In one embodiment, the replication application performs the quiescing function by sending "an instruction to production application 104, requesting the application to temporarily quiesce (i.e., temporary enter an inactive state)" and the application notifies the replication application when it enters a quiescent state so that the event marker can be stored only after the quiesced state is known to have been achieved. (JA79, 13:40-48.) This is consistent with other embodiments in which a "configuration utility" or the "replication application" "***places*** the production application in a quiescent state." (JA78, 12:41-44.) In all embodiments, some component of the system instructs the production application to quiesce.

The requirement that an application be instructed to quiesce is also consistent with the term's customary meaning in the art. (JA272 (defining "quiesce" as "***[t]o bring a device or an application to a state*** in which (a.) it is able to operate, (b.) all of its data is consistent and stored on non-volatile storage, and (c.) processing has been suspended and there are no tasks in progress (i.e., all application tasks have either been completed or not started).").) As the definition makes clear, "quiesce" is an active term that requires some entity to "bring" the production application into a quiescent state.

In an effort to expand the scope of prior art, Zerto proposes to construe "quiesce" to cover states in which an application's write activity happens to be paused on its own. That assertion is inconsistent with the ordinary meaning of quiesce and the intrinsic record. The patents explain that if a backup is created without quiescing the application, "it is ***likely*** that the backup occurred when the entire application was not in a consistent state." (JA81, 18:34-37.) This passage recognizes that if the application is not quiesced, it is of course possible through serendipity that the application may have happened to be in a consistent state when the event marker was placed. The patents teach, however, that if application consistency is desired, the system does not rely on that unlikely event, and instead quiesces the application and awaits notification that the quiesced state is achieved before placing the event marker. (JA79, 13:40-48.)

Zerto's proposal is inconsistent with both the terms' customary meaning and the intrinsic evidence and should be rejected. EMC's proposal is consistent with both and should be adopted.

**E. "write journal"**

| | |
|---|---|
| **EMC**: "storage in which a series of write transactions are stored along with information that can be used to identify an order of the write transactions" | **Zerto**: "an area of storage in which a series of write transactions are stored in chronological order" |

The parties dispute whether a "write journal" must store write transactions in chronological order, as Zerto proposes, or whether the journal can store transactions in any manner so long as it includes information that can be used to identify the chronological order of the write transactions, as EMC proposes. Just like a deli counter "take a number" system can service customers chronologically without making the customers stand in chronological order, computer systems can store information identifying a chronological order of transactions without having to store the transactions "in chronological order." The patents illustrate in Figure 5 "one ***simplified*** depiction" of a write journal in which the write transactions are stored in chronological order. (JA80, 15:28.) The specification is clear, however, that this simplified depiction is solely for ease of illustration and that "the addressing scheme within the write journal 111 may be much more complex than shown in FIG. 5." (*Id.*, 15:39-41.) Zerto's proposal would limit "write journal" to the "simplified" illustration shown despite the patents' teachings that the addressing

scheme "may be much more complex." EMC's proposal should be adopted because it is consistent with the patents' teachings that other (e.g., "more complex") addressing schemes can be used.

**F. "client file index"**

| | |
|---|---|
| **EMC**: "an index of client files" | **Zerto**: "an index of the backed up data items" |

Claims 3 and 6 of the '091 patent each recites a "client file index configured to maintain a record of the surface copy." In an exemplary embodiment, a client-file index is referred to as "an index of the backed up data items." (JA76, 7:15-16). Zerto seeks to blindly incorporate the specification language into the claims. Doing so introduces a term "***the*** backed up data items," that lacks antecedent basis in the claims, and is duplicative at best (and inconsistent at worst), with language already in the claims.

Zerto's construction seeks to construe the label "client file index" in a manner that specifies what the index maintains, but the claims already ***explicitly recite*** what it is that the client file index maintains, i.e., "a record of the surface copy." Zerto's construction should be rejected because the jury will likely be confused about how the "backed up data items" relate to the "record of the surface copy" that already appears in the claims.

The claim language "client file" is a label for the index. There is no need to make this straight-forward limitation complex. EMC's proposed construction is easy for the jury to understand and explains that the index catalogues client files, which may include the record of the surface copy recited in the claims.

## IV. U.S. Patent No. 7,577,867

The '867 patent discloses a system for consistent data recovery that inserts a "tag" into a journal of write transaction when one or more production applications are quiesced. (JA52, 7:27-38.) The system organizes its components into one or more "consistency groups" that each includes one or more "host device(s)," "data protection appliance(s)," and "logical unit(s)." (JA44, Fig. 8; JA51, 5:11-17.) A "data protection appliance manager" issues a request to the one or more data protection appliance(s) to cause applications running on the one or more host

device(s) to enter an inactive state. (JA52, 7:29-34.) While the applications are inactive, a "tag" is inserted into a journal to identify the point in time when the applications were in a known inactive state. (*Id.*, 7:34-38.) The applications are then allowed to re-enter an active state. A "replication facility" can later use the tag in the journal to recover the data as it existed when the applications were quiesced. (JA51, 6:21-24.)

**A. "data protection appliance (DPA) manager"**

| | |
|---|---|
| **EMC**: Plain and ordinary meaning. If the Court further construes this term, EMC proposes: *"a computer or a cluster of computers that issues requests to at least one DPA*." | **Zerto**: "a DPA that receives a cross-tag instruction, sends quiescence and tag commands, receives confirmations, and sends un-quiesce commands" |

The parties dispute whether a "data protection appliance manager" must be a data protection appliance ("DPA").[6] Essentially, Zerto reads "DPA manager" to mean "a DPA that manages," while EMC reads it to mean "the manager of a DPA." As a matter of grammar, EMC's reading is correct. No one reads "payroll manager" to mean a "payroll that manages." The customary meaning of the term supports EMC's proposed construction.

While the specification describes an embodiment in which the DPA manager is also a DPA (JA58, 20:54-56), the specification explains that embodiments are "illustrative" rather than "restrictive" (JA61, 25:3-10). Moreover, the inventors did not express a clear intent to limit the term to that particular configuration. In fact, claim 31 demonstrates that the inventors intended a DPA manager to be a component that could be separate from a DPA. Dependent claim 31 limits the system of independent claim 19 to a system "***wherein the DPA manager is one of the DPAs***." If, as Zerto proposes, the DPA manager had to be a DPA, claim 31 would be superfluous.

## V. CONCLUSION

For the above reasons, EMC respectfully requests that the Court adopt EMC's proposed claim constructions.

[6] The parties agree that a DPA is, *inter alia*, "a computer or a cluster of computers." (D.I. 86 at 2.)

MORRIS, NICHOLS, ARSHT & TUNNELL LLP

*/s/ Michael J. Flynn*

Jack B. Blumenfeld (#1014)
Michael J. Flynn (#5333)
1201 North Market Street
P.O. Box 1347
Wilmington, DE 19899
(302) 658-9200
jblumenfeld@mnat.com
mflynn@mnat.com

*Attorneys for Plaintiffs*

OF COUNSEL:

Michael N. Rader
Gregory F. Corbett
Allen S. Rugg
Nathan R. Speed
Elizabeth A. DiMarco
WOLF, GREENFIELD & SACKS, P.C.
600 Atlantic Avenue
Boston, MA 02210-2206
(617) 646-8000

Paul T. Dacier
Krishnendu Gupta
William R. Clark
EMC CORPORATION
176 South Street
Hopkinton, MA 01748

April 11, 2014

**CERTIFICATE OF SERVICE**

I hereby certify that on April 11, 2014, I caused the foregoing to be electronically filed with the Clerk of the Court using CM/ECF, which will send notification of such filing to all registered participants.

I further certify that I caused copies of the foregoing document to be served on April 11, 2014, upon the following in the manner indicated:

Adam W. Poff, Esquire — *VIA ELECTRONIC MAIL*
Monté T. Squire, Esquire
YOUNG CONAWAY STARGATT & TAYLOR, LLP
1000 North King Street
Wilmington, DE 19801
*Attorneys for Zerto, Inc.*

Matthew B. Lowrie, Esquire — *VIA ELECTRONIC MAIL*
Aaron W. Moore, Esquire
FOLEY & LARDNER LLP
111 Huntington Avenue, Suite 2600
Boston, MA 02199-7610
*Attorneys for Zerto, Inc.*

*/s/ Michael J. Flynn*

Michael J. Flynn (#5333)